Institute for Scientific Computing Research

Students and Faculty

Center for Applied
Scientific Computing

Developing New Scalable Solver Methods

Mark Adams

University of California, Berkeley

Abstract

Mark Adams continued to develop scalable solver methods for large-scale unstructured finite element problems in solid mechanics. He originally implemented these methods in the course of his research towards his doctoral dissertation at UC Berkeley. During the summer of 1998, he worked with the Linear Solvers Group in the Center for Applied Scientific Computing (CASC) at Lawrence Livermore National Laboratory (LLNL). His research at CASC focused on scalable solver efforts with regard to the software design of parallel linear solver libraries, and construction of scalable numerical primitives, common to his solver and that of the algebraic multigrid group in CASC. Mr. Adams also continued to discuss issues related to the design and use of PETSc (from Argonne National Laboratory) for the multigrid codes in CASC.

He continued to develop solution algorithms for contact problems with Lagrange multipliers, via Uzawa type algorithms, for the ALE3D group at LLNL. Mr. Adams has a serial prototype within his parallel code (using perfectly scalable constructs) that effectively solves some of the simple contact problems from the ALE3D group.

At the conclusion of the summer, Mr. Adams was able to solve problems exhibiting up to 9.6 million degrees of freedom in large-deformation, incompressible elasticity (Poisson's ratio of 0.49) and plasticity problems, with large jumps in material coefficients (1.e-4), on a Cray T3E and an IBM PowerPC cluster.

Change Management in the Digital Library Environment

Igg Adiwijaya

Rutgers University, Newark, New Jersey

Abstract

Igg Adiwijaya worked with the DataFoundry project, specifically in the area of change management of autonomous data sources. DataFoundry is a data warehousing and integration research effort underway in the Lawrence Livermore National Laboratory Center for Applied Scierntific Computing (CASC). Mr. Adiwijaya familiarized himself with DataFoundry and the sources from which the project extracts data. He also researched the methods for detecting and extracting changes to external data sources. He then eavluated the available approaches, including their advantages, limitations and potential for applicability to DataFoundry. He also implemented an automated process for extracting data from external data sources, and scheduling periodic runs.

Mr. Adiwijaya proposed a general framework for change management in the Digital Library environment, particularly the one serving DataFoundry. For DataFoundry, the framework requires approaches such as comparing two semi-structured documents, ranking semi-structured documents based on schema, and extracting changes to semi-structured documents. He devised an effective and efficient approach for meeting such requirements.

While concluding his summer with ISCR, Mr. Adiwijaya conducted further research on representation of changes and data mining on the extracted changes. He also started the implementation of the parser and delta extractor, components in the framework, and prepared one technical manuscript documenting the change management framework.

First Order System, Least Squares (FOSLS) Methods for Radiation and Neutron Transport

Travis Austin

University of Colorado, Boulder

Abstract

Over the summer of 1998, we began to implement an idea for accurately describing neutron transport within different types of media. Many current ideas are base on methods that suffer in particular types of media. For example, researchers currently are forced to deal independently with problems defined in the types of media exhibiting the property that neutrons have a small mean free path before collisions occur. We desire to treat such problems with the same ideas we are using in other cases. More importantly, we want our method to be successful, independent of problem characteristics.

In designing the problem formulation, we insisted that our solver have certain characteristics. This drove us to pose the problem as a least squares minimization problem. We coupled this idea with an expansion of our solution in spherical harmonics and a discretization of the moments with finite elements. To accurately treat all cases, we chose a scaling of the matrix equation based on the problem characteristics. Since we want a robust solver, we considered using multilevel techniques that have the potential to achieve convergence factors that are independent of mesh size.

The backbone of the described transport code has been established. Yet, we discover that our desire to use a multilevel solver and be robust compels us to pay special attention to solving for the first order moments. Multilevel solvers are extremely efficient and accurate when applied to certain classes of problem. Unfortunately, our formulation creates a system of equations for the first order moments that standard multilevel solvers do not handle well. This requires treating this case individually.

The second half of my summer dealt exclusively with this addressing this problem. We wanted to show that we could get good convergence results for the first order moments. With this established, we would be able to move on confidently to establishing the success of our method for the more general neutron transport problem. These characteristics were evident in our code, and this demonstrated that handling the first order moments as independent of the other moments was necessary. The code designed also showed that we could be successful in this endeavor.

Examining the Literature on Colloid and Colloid-Facilitated Transport Modeling

Lora Ballinger

University of Utah

Abstract

Current literature on colloid transport and colloid-facilitated transport was examined to evaluate the state of knowledge and the strengths and weaknesses of available models. Among the journals included in this literature examination were Colloids and Surfaces, the Journal of Contaminant Hydrology, the Journal of Nuclear Science and Technology, Nuclear Technology, Radiochimica Acta, Radiochemistry, and Water Resources Research. This literature search also included articles from conference proceedings (Scientific Basis for Nuclear Waste Management Symposia, Materials Research Society Symposia), and technical reports from LLNL and the Center for Nuclear Waste Regulatory Analysis. An annotated bibliography of 51 entries, which analyzes the relevance of the literature examined, was produced.

In addition, the relevant points from the literature were detailed in a report. In summary, two types of colloid must be considered: pseudocolloids, radionuclides bound onto natural colloids; and radiocolloids, composed entirely of radionuclides. Models are needed to describe the formation of each of these types. The typical colloidal transport model requires knowledge of parameter values that are not easily measured. In addition, in some cases, the theoretical predictions are insufficient to describe observed behavior; therefore, models based on this theory will not be accurate. A colloid velocity model must be developed that takes into account hydrodynamic chromatography effects. In addition, filtration, particle size and particle chemistry confer major effects on the distance traveled and the velocity of colloidal particles; these effects should be explicitly included in the model formulation. Finally, very few of the models in the literature have been validated adequately, particularly in field situations.

Based on information in the literature, recommendations were made regarding the necessity and feasibility of incorporating a colloid- facilitated transport model into the reactive transport model of the Yucca Mountain system. In some cases, the literature indicates that colloids could possibly have a significant effect at the Yucca Mountain site and should be explicitly included in the model. Most studies, however, consider colloid transport in a porous medium and do not consider the case of episodic flow in a fracture. Therefore, one cannot predict with any accuracy the effect of the colloids without including them in the model. This problem is further compounded by the fact that a majority of the data needed to predict parameter values for existing models is not available — either at all or for this site specifically.

It was recommended that an effort be made to obtain information on colloid diffusion coefficients and filtration constants (may involve surface potentials of the fractures, as well as density of binding sites and electrolyte concentration). Recommendations were also made for specific components of the model: As few colloid formation models exist, it was suggested that a simple equilibrium model could best serve for the time being. Particular functional forms for adsorption onto and desorption from colloidal particles were given. A simple means for combining two of the typical models for colloid velocity was presented. A basic form of the general transport equation was given, and a list of significant factors to was given. Among these significant factors are the distribution of fracture apertures and particle sizes, and the density of binding sites on a particular material.

Topics of interest in developing a model, including pseudo-colloid formation, filtration effects, and particle size and chemistry effects, were addressed. The availability of model parameter values at the Yucca Mountain site was determined by a literature search. Finally, a simple model of colloid transport was developed and implemented using the available parameter values. The model used a finite differences formulation to describe movement (not formation) of colloids within a 1D column. Results from the model runs indicate that, using the values gleaned from the literature, it would take only a couple of years for a particular location on the fracture to saturate, once the colloids arrived there. In addition, it was seen that by 10 years, the walls of a fracture 40 meters long would be completely covered. These could be good indications, or bad. If the number of binding sites were unlimited, quick binding would be good; radionuclides would be trapped on the walls of the fracture. If the number of binding sites were small, and the colloids could bind to each other, as well as to the fracture walls, then they could plug the fracture and prevent further transport. The results otherwise indicated, however, that the particles move very quickly in comparison to dissolved particles.

Research Progress in Pattern Classification, and Seminars Given on Object Recognition, Neural Networks, and Genetic Algorithms

George Bebis

University of Nevada, Reno

Visiting Faculty Abstract

During my visit to LLNL and CASC (July 8, 1998 through August 20, 1998), my major responsibility was to assist Sapphire Project members in identifying the preprocessing and pattern classification techniques appropriate for the FIRST image data (from the survey of Faint Images of the Radio Sky at Twenty centimeters). We concentrated on the problem of identifying radio sources exhibiting the special morphology known as "bent-double." This problem is quite challenging, as astronomers have a hard time defining bent-double morphology precisely. They also experience difficulty in confidently classifying a radio source image as a bent-double. I believe that many researchers find it very hard to come up with a general and complete definition for characterizing bent-double morphologies. This problem is thus very suitable for pattern classification techniques that extract knowledge about the problem at hand by examining a large number of examples (e.g., neural networks, decision trees, and the like). For these techniques to be successful, however, it is vital that we provide researchers with important information (features) extracted from the examples. This requires the use of sophisticated noise reduction and segmentation techniques.

In addition to my Sapphire Project responsibilities, I also gave one formal and two informal seminars. In the first, I reviewed the problem of object recognition, and then presented my recent work on the application of algebraic functions of views for object recognition. In the second seminar, I presented the fundamentals of neural networks, and then discussed an interesting application involving the use of neural networks for face detection from gray-scale images. In my final talk, I presented the fundamentals of genetic algorithms and discussed my recent work on the use of genetic algorithms for object recognition.

Work Completed

During the first phase of my work, I focused on the problem of feature extraction. My goal was to develop approaches for extracting features in a way that will allow pattern classification techniques to separate bent-double from non-bent–double morphologies. My activities are summarized in the following bullet points.

- **Appropriate image processing, neural networks, and machine learning software.** I did an extensive search to identify appropriate software to be used in the project. I also helped install and test some of the software packages.
- **Astronomical Image Processing System (AIPS).** AIPS is a popular software package for processing astronomical images. As I had no familiarity with this package, I spent some time trying to understand which of its functions might be useful in extracting important features from the radio source images. I also spent some time trying to understand FITS, the format in which astronomical images are stored.
- **Information stored in the catalog of the FIRST survey.** Astronomers who performed the FIRST survey have already extracted and catalogued some of these images. I spent a good amount of time trying to understand what information had been extracted from the images and how it can be useful for pattern classification. My conclusion is that the information stored in the catalog will be insufficient in helping to classify a radio source as a bent-double or alternative morphology. Additional features should be extracted (see Feature Extraction below). The information stored in the catalog, however, will help us identify promising regions of the sky, that is, regions where it very likely that radio sources can be found. Concentrating on these regions reduces the time needed to search the whole sky.
- **Reviewing the literature.** I spent a good amount of my time reviewing the literature. I found and studied many interesting papers on the application of classical pattern classification, neural networks, and machine learning techniques in astronomy, as well as in other related areas. Neural networks and

machine learning techniques generally have been applied to the problems of galaxy and star classification. Good results have been reported in these cases, which lends me to believe that they should also work well for our problem. Copies of these papers were provided to other Sapphire investigators.

- **Noise reduction and segmentation.** Before features can be extracted from the images, it is important to filter out the noise, apply some kind of image enhancement, and then separate the regions of interest from the background. During the limited period of time I stayed at LLNL, I experimented with several noise reduction and image enhancement techniques, and I came up with a procedure that seems promising. I implemented and successfully tested this approach. Chandrika Kamath and the other members of the group have a good idea about how this approach works.
- **Feature extraction.** After the regions of interest have been isolated, the next step is to compute features within those regions that might be useful for classification. Since it was not obvious to me which features were the most important for our problem, I decided to extract as many features as possible. In my current implementation, the features I extract are based on geometric properties of the regions (area, compactness) as well as histogram features (texture, moments).
- **Feature selection.** Presenting all possible features to the pattern classifiers would very likely overwhelm them, particularly as some of these features carry no discrimination power benefiting the problem we are trying to solve. This technique would yield very poor results. Thus, it is essential to first apply feature selection—techniques that reduce the dimensionality of features before they are fed to the classifier. A popular method known as principal components analysis (PCA) considers linear combinations of features. I implemented PCA, but did not have time to test it during my stay at LLNL.
- **Classification.** I designed a method for distinguishing bent-double from non- bent–double morphologies based on neural networks. Due to my limited visit, however, I did not have the time to implement and test it. I discussed this method with the other members of the Sapphire Project, as well as with Bob Becker from the LLNL Institute of Geophysics & Planetary Physics (IGPP). He finds it interesting and believes it is worth trying.

Work to be completed

- **Parameters for noise reduction and segmentation.** The procedure I have implemented for noise reduction and segmentation can give very good results but it depends on some parameters. If these parameters are not chosen appropriately, the segmentation results can be very poor. I implemented a number of heuristics to choose these parameters adaptively; however, more heuristics and much more testing are needed in order to come up with a scheme delivering very good performance. I should emphasize that the noise reduction and segmentation steps are very important and will influence the performance of the whole system if the results they produce are not very good.
- **Feature extraction, selection, and normalization.** We should consider extracting additional features that can provide more information to the classifiers. Additionally, the PCA approach as well as other approaches must be applied for selecting only the most important of the features we intend to feed to the classifiers. Issues related to the normalization of the features must also be considered, especially when neural networks are used for classification.
- **Non-bent–double training examples.** Bob Becker and his group have already provided us with images of radio sources that likely correspond to bent-doubles. However, in order to train a classifier to distinguish bent-doubles from non-bent–doubles, a representative number of non-bent–double examples must be included in the training set. Careful selection of the non- bent–double examples is critical in minimizing false positives as well as false negatives.
- **Appropriate classification schemes.** A large variety of classification techniques is available for investigators to try on a given problem. I believe that after we have come up with a good set of features, it becomes important to test various neural network and machine learning techniques. Based on my experience, finding the best technique for a problem as well as fine- tuning it (e.g., choose the appropriate number of hidden nodes in a neural network) requires a lot of experimentation. Recently, many researchers are interested in combining the results of several different classification techniques (fusion) in order to improve performance. Results based on this idea have been promising, and I believe that we should also consider this direction.

Wavelet-Based Compression Schemes

Martin Bertram

University of California, Davis

Abstract

I have implemented a wavelet-based compression scheme for time-varying volumetric data sets. For the wavelet transform, wavelet lifting, in-place computation of the coefficients and integer arithmetic were applied. After several experiments it turned out that a linear B-Spline wavelet performed best in decorrelating the data for compression.

I also developed an efficient coding scheme. This scheme can be further improved by taking into account the hierarchical structure of the wavelet transform and different probability models for each single coefficient or each bit of a coefficient.

High-Performance Surface Rendering

Kathleen S. Bonnell

California State University, Dominguez Hills

Abstract

I worked with Dan Schikore on high-performance surface rendering. In particular, I investigated performance differences between different triangular models. My first tests involved the SGI Infinite Reality Engine benchmark program, which renders 10 million triangles per second using a triangle strip representation of the data set. Currently-render performance of individual triangle models yields only about 10% of that stated performance of the architecture. For my tests, I measured how varying the length of the triangle strips affected render performance. The motivation behind these tests was to increase performance. Individual triangles have a cost of 3 vertices per triangle, whereas triangle strips have a number of triangles plus 2. It would seem to follow that the longer the strip one could create, the better the performance should be.

My tests indicated otherwise. In fact, performance peaked at a range of 10–12 triangles per strip; after that it leveled off, and longer strips did not help to improve performance. On the flip side, the test results dramatically illustrated the fact that strips of length less than 10 were a definite hindrance to performance.

After establishing this, I moved on to tests involving generalized triangle strips (GTS). Conceptually, generalized triangle strips can turn corners. Thus, one can generate extremely long strips by turning enough corners. Old GL hardware supported GTS via a swap facility with minimal cost. Newer GL hardware can only simulate a swap by sending a duplicate vertex to the renderer, with an obvious cost of one vertex per swap. We wanted to test the performance of strips with swaps as opposed to strips without swaps on the same data set. To that end, I modified an existing software package (Stripe by Francine Evans, State University of New York at Stony Brook). I modified the package to work with isosurfaces and material boundaries for ASCI data. I also created a no-swaps version of the software. Results of performance tests with this modified software package demonstrate that swaps performed better than no-swaps; and minimizing swaps performed better than not minimizing.

At first these results might be surprising. The Stripe package, however, was designed to create the longest strips possible, regardless of the number of swaps involved. When I removed the swaps from the strips, this created a great number of strips of length less than 10. Based on the tests with the SGI benchmark, the poor performance of my no-swaps version is understandable.

The next step would be to test an algorithm that creates strips with a length of at least 10 triangles— both a swaps and a no-swaps version. Unfortunately, I ran out of time, so that must be left for a future project.

Algebraic Multigrid Visualization

Tim Chartier

University of Colorado, Boulder

Abstract

For twelve weeks during the summer of 1998, I worked as a summer student at Lawrence Livermore National Lab (LLNL). In particular, I worked within the Center for Applied Scientific Computing (CASC). Van Henson and Jim Jones guided my research.

A focus of the multigrid research team at LLNL is to apply algebraic multigrid (AMG) technology to a host of problems of importance to the Lab. CASC intends to further develop AMG and to find parallel algorithms.

My summer began with work on AMG visualization. To aid in research, it is often helpful to see the grids that AMG produces and to plot the residual for the cycles. Using Matlab, I produced pop-up windows to visualize both of these important aspects of AMG. Others and myself used this tool throughout the summer. The tool was designed for my work in AMGe, which is a new version AMG based on element interpolation. This tool can easily be transported into existing AMG code used at LLNL.

A goal of the multigrid research team at LLNL is to numerically solve elasticity problems. AMGe shows promising results in such problems. Yet, an open question is how to find "good" grids for AMGe. In two-level multigrid, we want to solve for the fine grid error on a coarse grid, and then interpolate this error back to the fine grid. My goal was to begin research on such coarsening.

AMGe offers new tools for deciding what sets of points would give good interpolation. These tools supply large amounts of new information. How could this information help us find numerically useful grids? With the help of Jim Jones and Van Henson, I implemented an algorithm to begin to answer this question.

As the summer progressed, we gained valuable insight on our new AMGe tools. Useful discussion brought new results. These results further developed options and ideas for future research. By the end of summer 1998, I left LLNL with many ideas for further research in coarse grid selection for AMGe. AMGe shows promise as a numerical solver for difficult and important problems such as elasticity. Such an algorithm would help LLNL with its projects and applications. AMGe exhibits promise both in its usefulness and its robustness.

Rayleigh-Taylor Mixing: Computational Investigations of Linear Electric Motor Experimental Data

Jonathan Dursi

University of Chicago ASCI Alliance[1]
Tri-Lab Alliance Technical Sponsor: Stephen Libby

Abstract

At the heart of understanding the accuracy of interface mixing simulations lies the careful validation of numerical models with correctly interpreted experimental data. In this study, Jonathan Dursi performed a variety of geometrical and statistical studies of data from linear electric motor (LEM) fluid interface experiments by Dimonte et al. at Lawrence Livermore National Laboratory). Such studies are a necessary prerequisite to any serious experiment–theory comparison.

The LEM data consists of 2D slices through a 3D Rayleigh-Taylor mixing zone, with two immiscible fluids of Atwood number of 0.34 and subject to an acceleration of 75 G. In each experiment, images from three time frames show the deeply non-linear evolution of the mixing zone. To understand the mixing physics, the approach—rather than to examine various spectra of the resulting density field—was to investigate individual "blobs" for spikes of heavy fluid penetrating the lighter fluid, and "bubbles" for light fluid intrusions into the heavier fluid. The images from the experiments had already been processed (smoothed and bi-leveled) into values which represented either the high- or low-density fluid. The blobs could then be easily identified and tabulated by performing a linear search through the images, using a flood-fill algorithm to identify connected regions. This allowed some very simple investigations to be performed.

Dursi began by examining the evolution of the number of blobs as a function of the penetration depth. Although many mixing models typically use very peaked distributions of blobs, the distribution was clearly seen to quickly become quite flat. Similarly, he investigated the size distribution of the blobs, and the evolution of the distribution. The distribution again seemed remarkably flat, but with a strong peak for small blobs; however, "slicing" effects through the 3D experiment will tend to bias results toward small, observed blobs. It is not yet clear to what extent this bias affected the results.

More sophisticated analysis is underway, involving modeling the observed blobs with simple shapes, and noting the resulting size/shape distribution over time under constant and varying accelerations. This analysis, as above, requires untangling the effects of taking a 2D slice through 3D objects. Further work on this project involves comparing experiments with simulations, and examining more carefully the image processing procedure to determine sensitivity of results. When 3D simulation results are available, the experimental procedure for image acquisition and processing will be followed within the simulation cube, to both help understand the experimental data, and to potentially verify the simulation results. Sensitivity analysis will be done by examining sensitivity of results to varying the procedure and parameters of the image post-processing.

[1] The University of Chicago ASCI Alliance is focussed on advancing the understanding and modeling of astrophysical thermonuclear flashes that are due to the accretion of material onto a compact body and subsequent ignition. These flashes manifest themselves respectively as x-ray bursts, type 1a supernovae, or novae depending on whether the accretion is onto a neutron star or a white dwarf (with either an off-center carbon flash for a supernova 1a or a thin shell hydrogen flash for a nova). These flashes are significant in the production of light elements, in the accurate determination of the cosmic distance scales, and in probing the physics of the compact objects themselves. Since the accumulation of material and its ignition mode depends crucially on turbulent mixing, a significant part of the effort at Chicago is focussed in this area. In particular, the effort of the two LLNL summer students, Jonathan Dursi and Yuan-Nan Young, was aimed at advancing and validating Rayleigh-Taylor interface mixing models, layer formation, and subgrid turbulence models.

Multigrid for Problems with Local Refinement

Christopher Higginson

University of Colorado, Boulder
Department of Applied Mathematics

Abstract

To obtain efficient multigrid solvers for problems with discontinuous coefficients, generally one must take the problem coefficients in to consideration in defining the components of the multigrid algorithm. For example, in the so-called "black box" multigrid method, the operators used to transfer functions between grids are defined based on the coefficients of the fine grid matrix. at the Center for Applied Scientific Computing (CASC) at Lawrence Livermore National Laboratory, these black-box operator-based interpolation formulas have been used successfully in two parallel multigrid solvers, SMG and ParFlow.

Applying these operator-based interpolation formulas when the global grid is logically rectangular is relatively straightforward. Together with Jim Jones of CASC, Chris spent the summer looking at methods for applying these black-box multigrid ideas to grids with local refinement, where the global grid will no longer be logically rectangular. In particular, he looked at refinement for cell-centered discretizations.

Chris looked in the literature for existing robust cell-based multigrid solvers. He implemented several of these and tested them on difficult diffusion problems with discontinuous coefficients. Based on their lack of robustness, Chris developed a new algorithm for cell-based multigrid. By careful construction of an interpolation operator that took into account variation in the diffusion coefficients, he was able to improve multigrid performance. The new algorithm was successful on several test problems where the existing methods had failed. Together with Jim Jones, Chris wrote a draft paper describing the new method. This work should provide some valuable information about how to transfer residuals and function values at the fine-coarse interface, an issue that must be addressed as researchers build general purpose multigrid solvers for use in codes with structured adaptive mesh refinement.

Chris also spent part of the summer looking at a first order system, least squares (FOSLS) approach to parabolic problems. This approach has the potential to provide rigorous error estimates to guide refinement in both space and time. Chris developed the discretization method and began looking at parallel space-time multigrid solvers for the resulting discrete equations.

A Problem-Solving Environment for the LOCFES Suite

Rithea Hong

Texas A & M University

Abstract

My project for the summer, under the supervision of James McGraw, was to design and implement a problem-solving environment for the LOCFES suite. This suite is a group of one-dimensional neutron transport codes that my thesis advisor, Dr. Paul Nelson, has developed over several years. These codes solve the monoenergetic, steady state, one-dimensional neutron transport problem in plane parallel geometry. They are tools designed to aid in the study of various spatial discretization schemes. The main elements of my project were:
Identify essential elements of an environment that would make using these codes easier and more efficient.
Design the environment.
Implement the environment.

At the beginning of the summer, I had only a conceptual idea of what the environment should contain. My work began by identifying what would make the user's life easier and more efficient. The existing suite of codes operates by reading in an input file containing all of the relevant parameters, and then, eventually, producing an output file containing the results. Although it is simple to code, such an interface is a barrier to the user. The user must first learn the details of the syntax for the input file before using the code. Also, once the user has learned the code, he or she must handle the management of input and output files individually. All this is tedious to the user and wastes time. I designed an environment that would provide for interactive problem specification, result analysis, and management of both input parameters and results. During the design phase, I obtained feedback from several people, including those familiar with neutron transport and those familiar with GUI design. Their input proved extremely useful.

Before implementation began, it was necessary to survey the available development environments and to weigh carefully each of their strengths and weaknesses in terms of what the project required. I felt that the most desirable development environment would be sufficiently powerful to implement the designed system, be sufficiently mature so that defects in the system would not interfere with the implementation, be accessible so I could actually use it, and be widely used. The final requirement ensures that an infrastructure exists to aid during development, and later maintenance. One does not want to choose an orphaned environment. After weighing several options, I concluded that Java was the best option. Since it was designed in the era of GUIs, it provides powerful support for the rapid development of the environment I envisioned.

After selecting Java, I spent some time familiarizing myself with the language before beginning implementation. After becoming comfortable with Java, I began implementing the GUI portion of the environment. By the end of the summer, the GUI portion was largely complete. Remaining tasks include: implementing the storage and retrieval of parameters and results, integration with the LOCFES code, and analysis of the results.

During this work, I also attended numerous presentations and interacted with a wide variety of Livermore personnel. I thoroughly enjoyed the experience and made considerable progress in my project. Particular thanks go to James McGraw, Neale Smith, Neil O'Neill, and Milo Dorr who helped me with my projected throughout the summer.

Amino Acid Descriptor Clustering Tools

Rachel Karchin

Stanford University

Abstract

During the summer of 1998, Rachel Karchin participated in the Lawrence Livermore National Laboratory summer student program by working with researchers in the Center for Applied Scientific Computing (CASC). She wrote Perl scripts to implement a web page that queries the DataFoundry warehouse. The web page visitor selects from a wide variety of attributes and conditions available on the page, and the scripts produce SQL query statements, execute them, retrieve data from the server, and display the results in tabular format.

Karchin also created several new html pages and Perl scripts to implement amino acid descriptor clustering, using the DataFoundry warehouse. These tools enable a user to:

- View all the amino acid descriptors in selected pdb chains.
- Perform computations on descriptor columns.
- Create new clusters of amino acid descriptors by selecting and launching a clustering program.
- View clustering results and write them back to the warehouse with additional annotation, the date of the experiment and a description, entered by the user.
- View the most probable sequence within a user-selectable window size for an existing cluster of descriptors.

Additionally, Karchin researched available machine learning public domain software for the Sapphire project and installed and documented the decision tree programs oc1, ripper, and cn2 and the clustering program autoclass. Independently, she studied C++ and machine learning; took several Lab TV courses, including Java Projects, C++ iostream, and Calculus and Probability Refresher for Engineers; and attended object-oriented design users group meetings.

Before leaving CASC, Karchin gave a 45-minute presentation on hidden Markov models and their applications to protein classification and multiple alignment. She presented this material to a group of Lab Biology and Biotechnology Research Program (BBRP) scientists.

Karchin demonstrated the SAM software suite to a group from BBRP and CASC's DataFoundry Project. SAM was written at UC Santa Cruz, where Karchin was one of the coders. It implements hidden Markov model technology for use with protein, DNA and RNA sequences. She also installed a local copy of SAM on one of BBRP's SGI servers.

Visualization of Large Scientific Data Sets Using Virtual Reality

Falko Kuester

Center for Image Processing and Integrated Computing (CIPIC)
University of California, Davis

email:kuester@cs.ucdavis.edu
graphics.cs.ucdavis.edu/~fkuester

Abstract

This summer research project, conducted at Lawrence Livermore National Laboratory (LLNL), in the Center for Applied Scientific Computing, focused on two distinct topics. The first was real-time visualization of large scientific data sets, using optimally adapting meshes and virtual reality (VR) display technology, The second was the motivations and benefits for the use of VR as a possible next-generation visualization and modeling environment. The requirement and implementation ideas for adaptive meshes are noted below.

Interactive Visualization of Large Data Sets

The key idea of this project was to identify an appropriate modeling scheme necessary for the visualization of huge complex scientific data. The Real-time Optimally Adaptive Meshes (ROAM) algorithm developed by Mark Duchaineau and others was selected as the underlying design idea and possible improvements were investigated. Our primary interest was to identify a real-time–capable surface or volume model satisfying the specific performance requirements listed below. We considered a variety of models, including parametric, implicit and polygonal representations.

Requirements

The following performance and feature requirements were established

- General surfaces (e.g. ISO-surfaces).
- General viewpoint notion.
- Highly agile view direction.
- Parallel run-time optimization.
- Improved run-time efficiency.
- Incremental frustrum culling.
- Visibility culling.
- Time dependent surfaces.
- Topological changes vs. scale.
- Texture hierarchies for general surfaces.
- Performance comparison to estimated true optimum.
- Better performance model.
- Geometry paging from disk (preferably with compression).
- Texture paging from disk (preferably with compression).
- Non-pre–computed adaptation of meshes.
- Multi-pipe display.

Related issues included hardware model for performance optimization, modifiable display lists, and parameter controls.

Results

The evaluation of surface models indicated that implicit surfaces and certain parametric representations like the Cao En surface model promised the best results. Demo applications have been created for these surface types.

Visualization using VR Technology

Problems associated with the visualization of the massive scientific data sets generated at LLNL demand new technologies. As part of this project, a feasibility study of virtual reality as a next generation visualization environment was launched. Preliminary results are discussed at the end of this section.

Terminology

In the early Seventies, Myron Krueger defined the term artificial reality, which was followed by derivatives such as cyberspace, virtual reality, virtual environments(VE) and synthetic environments (SE). The last term was

defined in 1992 by the Committee on Virtual Reality Research, established by the National Research Council. The terms VR and SE are used interchangeably in the following sections.

Vision

Push your keyboard to the side, put on your data gloves, turn on your 3D display and engage a truly 3D visualization and modeling environment. Either visualize pre-computed results or start from scratch by modeling the simulation environment itself. Beginning with geometric primitives or a pile of "virtual clay," and with the help of virtual design tools, create the initial model for a fluid mechanics or thermodynamics simulation. Communicate with other scientists through a centralized database-oriented design-space giving instant access to the latest 3D visualizations. Verify the setup and boundary conditions by running a few simulation cycles and visualize them in the created environment before spending hundreds of CPU hours on a problem with incorrectly defined boundary conditions. This mechanism would allow scientists to turn their insight into concepts, concepts into models, and finally models into production data or research results. The objective is to develop a suite of tools that will dramatically reduce cost, time, and complexity in research and development.

For the initial study, the focus was on the visualization aspects. The additional modeling component in this vision could be a powerful application area for the modeling schemes investigated.

Requirements

Synthetic 3D worlds require a completely different interaction and navigation paradigm than the classical keyboard-and-mouse approach. This paradigm might even be paired with appropriate visual and environmental clues, such as tactile feedback and 3D sound. Although 3D mice and pointers might seem appropriate for common VR applications, including architectural walk-throughs or terrain visualization, they are not suitable for the envisioned visualization and design tasks. Therefore, the feasibility study is targeting a dual-handed design paradigm in which the user is equipped with a pair of pinch gloves as the main input and interaction device, in addition to the previously mentioned 3D mice and pointers. The pinch gloves are equipped with basic gesture recognition, which allows the association of natural-hand movements with specific model interactions. Previously complex tasks, such as positioning, rotating, scaling, and modification of objects in 3D space become possible with intuitive hand gestures. This approach allows the preservation of the hands-on experience from the physical world, while overcoming the well known classical 2D constraints introduced with the keyboard.

The project utilizes an immersive workbench, a state-of-the-art VR projection device, that produces table-size, stereoscopic renderings of large-scale 3D data sets. This VR display device allows the projection of 3D computer-generated images onto an approximately 7-foot by 5-foot projection area. Its similarity in appearance to a drafting board brought it early recognition and acceptance by engineers. Even though the user is required to wear a special set of shutter glasses, this environment is semi-immersive. A user interacts with an object that is rendered stereoscopically in 3D space thus, providing the analog of the traditional 3D design space. The scientist can either use data gloves or 3D pointing devices to interact with a virtual model representing either the results of a simulation or its original problem definition.

The user, for example, will be able to grab two points in space and rotate either a selected object or the entire environment around the center point. A truly powerful feature of this implementation is its ability to provide an unprecedented amount of real estate to the user in the form of a 3D desktop. Anyone with exposure to working with multiple open and overlapping windows or virtual 2D desktops on a regular display will appreciate that objects, tools, and other components can now be placed, arranged and viewed in an unlimited 3D domain.

Results

A proof-of-concept framework supporting the mentioned input and output devices was designed and implemented in late summer 1998, and an initial visualization test on a demo data set from LLNL were performed. The user can now visualize an animation sequence of iso-surfaces derived from a 512x512x512 data set, and intuitively inspect it while immersed in a synthetic environment. More advanced support for user interaction was available in the fall.

Finite Element Modeling in Simulations of Full Flexion of the Human Finger

Jonathan Pearlman

University of California, Berkeley

Abstract

The bulk of my time this summer was spent running finite element analyses (FEA) on a model of a human index finger. The finite element model (FEM) previously developed in our Lab includes all of the bones of the finger and many of the soft tissues. Our primary goal this summer was to successfully simulate full flexion of the human finger, defined as a 21-mm excursion of the flexor digitorum superficialis (FDS) and flexor digitorum profundus (FDP) tendons. Our secondary goal was to validate the model by performing simulations of the index finger at postures described in the literature, and comparing joint reaction forces between results from the FEA experiments and values noted in the literature. Simulations addressing both of these goals were run in parallel throughout the five weeks that I worked at Lawrence Livermore National Laboratory.

Index finger flexion simulations were performed under minimal constraints: The metacarpal was fully constrained, with the three more distal bones unconstrained. A displacement of 21mm was prescribed for the most proximal elements of the tendons over a given time span. All simulations were done quasi-statically. Soft tissues (ligaments, annular pulleys, and tendons) were modeled as hyper-elastic material that was developed specifically for simulating this type of biologic tissue. A pre-stretch was prescribed for each of the ligaments, which, during the initial increments of the simulation, contracted all of the joints so that adjacent articular surfaces were in contact. Following this pre-stretch phase, tendon displacements resulted in flexion of each of the three joints of the finger. The success of a given simulation was judged by the amount of tendon displacement achieved before the simulation failed, and on how equally each of the joints flexed. (If one joint flexed much more than others it was clear that the simulation did not accurately model in-vivo finger flexion.) In the simulations performed this summer, the maximum tendon displacement was 14mm. All simulations failed due to soft tissue inversions (mostly in the tendons.)

To validate the FEM technique, certain postures (joint flexion angles) must be simulated. Using joint angles for a ŒTip Pinch, from Chao, et al., rigid body rotations were applied to the proximal phalanx, medial phalanx, and distal phalanx. There was limited success in these simulations primarily because of the difficulty of simultaneously displacing the tendons while imposing a rotation on the bones. Invariably, elements in the soft tissues would invert, causing the analysis to fail.

In an effort to improve the convergence of the analysis, the model was continually refined. One of the major refinements added a volar plate on the proximal end of the proximal phalanx. This addition was done to make the ligaments of the metacarpophalangeal (MCP) joint physiologically accurate. The dorsal ligaments of the MCP joint insert into the dorsal plate. Prior to the addition of the volar plate, the dorsal ligaments were attached directly to the proximal phalanx; this drastically reduces the moment arm the ligament has on the bone. The addition of the volar plate, and the subsequent re-meshing of the ligaments of the MCP joint, greatly improved the convergence of the model.

Reference: Chao, E.Y., K. An, W. P. Cooney, and R.L. Linscheid (1989), *Biomechanics of the Hand,* (World Scientific Publishing Co.)

Testing the ZPL code on ACI Blue–Pacific

Lyn Reid

University of Washington

Abstract

My summer project at Lawrence Livermore National Laboratory (LLNL) involved working on neutron transport code with Peter Brown and Anne Greenbaum. Peter is an LLNL researcher in the Center for Applied Scientific Computing (CASC), and Ann, like myself, was also visiting from the University of Washington for the summer. Under her supervision the previous year, I had developed code for one-dimensional steady state neutron transport in a parallel language called ZPL. Developed at the University of Washington, ZPL is a language that ports easily between sequential and parallel machines.

My project began by putting this ZPL code on the IBM ASCI Blue–Pacific machine, and comparing its performance with MPI. This testing uncovered some problems with the ZPL language that will be addressed by the developers. It also served to familiarize me with numerical methods for solving the transport equation and with parallel computing.

With this background, we then began to consider the problem of adaptive mesh refinement in transport codes. We discussed a number of ideas for adaptively refining the mesh or for using an irregular mesh created at the beginning of a run. When working in multidimensions, one must decide how to define the difference scheme at the interface between the underlying grid and a patched zone. Then, one must employ efficient methods for solving the matrix equations that come from such a difference scheme. It is best to retain the ability to march through the grid in the direction of neutron travel, and the difference scheme should allow for this. At the same time, one wishes to take advantage of parallelism as much as possible, and to generate matrix equations for which iterative methods converge rapidly. With many ideas to consider here, and we did not settle on a single strategy that seemed best. Work in this direction is continuing.

While I was at the Laboratory, I also attended several seminars sponsored by CASC. I had an enjoyable summer, and I feel that I learned a lot in my time spent at LLNL.

Visualization for SAMRAI Data Sets in AVS/Express

Michael Wittman

University of California, Davis

Abstract

Visualization of data created by adaptive mesh refinement (AMR) codes, such as the structured adaptive mesh refinement application infrastructure (SAMRAI) code being developed at Lawrence Livermore National Laboratory, normally is not supported by commercial visualization tools such as AVS/Express. While custom-built visualization tools exist for AMR grids, these do not have the flexibility and rapid development capabilities of AVS/Express. One facet of the SAMRAI Project explores the use of AVS/Express to see whether, with modest development work, it can be extended to meet data visualization needs.

One of the main problems in visualizing AMR data occurs at the boundary interfaces between multiple resolution patches. AVS/Express handles multiple patch data by running the visualization algorithm separately on each patch, and then rendering the individual patches as separate collections of data. This causes problems for techniques that use interpolation methods, such as isosurface and contour creation. Discontinuities are created along patch boundaries when coarse and fine patches are adjacent to each other. The focus of my work was to investigate solutions that reduce these discontinuities.

The first approach explored an unstructured representation and added sutures of additional AVS/Express cells to correctly connect the cells of each of the patches. This added some additional overhead but eliminated the discontinuities, while remaining very close to the original data. After some experimentation, it was found that this approach would require more time than was available to finish the implementation.

A second, and more traditional, approach was then investigated. Rather than adding cells to create continuity at the patch boundary, the data at the fine level were adjusted to match the course level data. This method allowed the standard AVS/Express modules to function on each patch separately, while creating a smooth interface for the isosurface module. Since this can cause some change in the isosurface created, users were consulted to verify that the change would not be significant to their understanding of the data. The second approach was completed and yielded acceptable results. The software developed will be integrated into the SAMRAI visualization toolkit.

Rayleigh–Taylor Mixing Instability in Miscible Fluids and Layer Formation in Stably Stratified Fluids

Yuan-Nan Young

University of Chicago ASCI Alliance[1]
Tri-Lab Alliance Technical Sponsor: Stephen Libby

Abstract

Rayleigh–Taylor Mixing Instability in Miscible Fluids

In this work, Yuan-Nan Young studied the Rayleigh–Taylor instability with particular attention to the mixing zone at the interface in miscible fluids. This research was conducted in collaboration with Guy Dimonte, Andrew Cook, and Oleg Schilling of Lawrence Livermore National Laboratory.

An important experimental feature herein addressed computationally was the asymptotic scaling of the mixing zone as _Agt2 (where __is a constant equal to about .05, A is the Atwood number, and g the acceleration). This study focussed both on the accuracy of the scaling law and the dependence of the constant __on the form of the initial perturbation. The model system studied was the incompressible Navier–Stokes equation in the Boussinesq approximation (the density contrast for the system causing the Rayleigh–Taylor instability is achieved through linear density dependence on a scalar field). In this case, Atwood numbers are naturally quite small, ~10^{-3}. To model this system, Young modified and adapted a 2D spectral code originally created by Cattaneo, Werne, and Julien. The calculations were done on a 5122 grid. Preliminary calculations for the idealized case of single modes found growth rates similar those found for immiscible viscous fluids. These studies also demonstrated strong dependence of the early mixing on perturbation wavelength: The higher the perturbation wavelength, the faster the mixing zone grew. Further investigation showed that the mixing zone followed the _Agt2 rule independent of initial conditions. However, _ was discovered to vary from .02 to .06, depending on initial conditions.

Layer Formation in Stably Stratified Fluids

Layer formation is thought to be an important process in astrophysical phenomena (e.g., semi-convection). A detailed understanding of energy transport and mixing of chemical compositions within layers is therefore highly desirable. Since, from a numerical perspective, it is not yet possible to simulate these phenomena directly, it is appropriate to attempt to capture the physics in a reduced model. One such model is the one-dimensional turbulent model (ODT) designed by A. Kerstein (Sandia National Laboratories/California), which subsequently has been applied to scalar field mixing by Kerstein and S. Wunsch (Chicago, and Sandia National Laboratories/California, respectively).

The problem of layer formation in stably stratified fluids is an example of a test bed in which ODT-based models can be compared both to experimental results and to direct numerical simulation. In this work, Young (in collaboration with Kerstein and Wunsch) took the approach of formulating a toy model to study layer formation in a stirred stratified fluid (a case first studied by Linden and Whitehead). He then applied ODT and direct numerical simulation (DNS) to the model, and compared the results with ODT. Preliminary ODT model and DNS calculations show that, upon adjustment of one free parameter, the ODT model yields agreement with the layer width obtained via DNS (using the same forcing in both cases).

[1] The University of Chicago ASCI Alliance is focussed on advancing the understanding and modeling of astrophysical thermonuclear flashes that are due to the accretion of material onto a compact body and subsequent ignition. These flashes manifest themselves respectively as x-ray bursts, type 1a supernovae, or

novae depending on whether the accretion is onto a neutron star or a white dwarf (with either an off center carbon flash for a supernova 1a or a thin shell hydrogen flash for a nova). These flashes are significant in the production of light elements, in the accurate determination of the cosmic distance scales, and in probing the physics of the compact objects themselves. Since the accumulation of material and its ignition mode depends crucially on turbulent mixing, a significant part of the effort at Chicago is focussed in this area. In particular, the effort of the two LLNL summer students, Jonathan Dursi and Yuan-Nan Young, was aimed at advancing and validating Rayleigh-Taylor interface mixing models, layer formation, and subgrid turbulence models.

ISCR Annual Report FY 1998

Lawrence Livermore National Laboratory
P.O. Box 808, L-561
Livermore, CA 94551